

METHOD OF HIERARCHICAL CACHING OF CONFIGURATION DATA HAVING
DATAFLOW PROCESSORS AND MODULES HAVING TWO- OR
MULTIDIMENSIONAL PROGRAMMABLE CELL STRUCTURE
(FPGAs, DPGAs, etc.)

Background information

Related art

The related art on which this patent specification is based is
5 described in Patent Application 196 54 846.2-53 (Method of
automatic dynamic reloading of dataflow processors (DFPs) and
modules having a two- or multidimensional programmable cell
matrix (FPGAs, DPGAs, etc.) and in Patent Application 196 54
593. 5-53 (Run-time reconfiguration method for programmable
10 modules). A method of configuring and reconfiguring DFPs, as
well as FPGAs, DPGAs and similar modules according to the
related art in which a separately configured central higher-
order microcontroller-like module assumes the task of
distribution of configuration data to a plurality of lower-
15 order, mostly passive control units is described in these
documents.

Disadvantages of the Related Art

20 By using a central global unit which controls the
reconfiguration of parts (e.g. cells (CELS)) of one or more
modules, bottlenecks occur when a plurality of different
reconfiguration requests are to be handled at the same time.
The advantages of the parallelism of the above-described
25 modules are considerably limited by such a central unit, since
it represents the typical bottleneck and substantially slows
down the processing of data.

Furthermore, assigning the event source to the configuration
30 to be loaded represents a problem because absolute addresses

of the configuration memory are used. The reconfiguration unit must therefore contain a type of memory management system which, like in an operating system, also documents which memory area is used by which configuration.

5

Management of resources (e.g. CELs) represents an additional problem. It must be ensured that each CEL is assigned exactly once to each algorithm started by a reconfiguration request and, specifically, to the one that also uses the remaining 10 surrounding CELs; otherwise deadlocks may occur.

In order to elucidate the problem of reconfiguration again, the following example is given: a matrix of CELs is 15 reconfigured and in the RESET state. Each CEL is capable of indicating whether it is in a reconfigurable state. All CELs in the matrix are ready to be configured; thus they are in a reconfigurable state. A first configuration routine (KR1) is loaded; the matrix is not fully utilized. The configured CELs clear the indication that they are in a configurable state. A 20 second configuration routine (KR2) independent of the first one is loaded in a group of not yet configured CELs. A third configuration cannot be loaded, since this requires CELs of the first and/or second configuration routine (KR3); however these are not in a reconfigurable state as they are being 25 used.

KR3 must be stopped until the required CELs are released, i.e., KR1 and KR2 are terminated.

30 During the execution of KR1 and KR2, a load request for a fourth configuration routine (KR4) and a fifth configuration routine (KR5) arrives, which cannot all be loaded immediately, because they use CELs that are being used by KR1 and KR2. KR3 and KR4 partially use the same CELs; KR5 uses none of the CELs 35 of KR3 and KR4.

In order to properly reload KR3 - KR5, the following

requirements must be met:

1. KR3 - KR5 should be loaded in the order of the load requests if possible.

5

2. As many KRs as possible that are independent of one another, i.e., have no common CELs, should be loaded in order to achieve maximum parallelism.

10 3. The KRs should not block one another, i.e., KR3 is partially loaded but cannot be loaded any further since other CELs are blocked by the partially loaded KR4; while KR4 also cannot be loaded further since again required CELs are blocked by KR3. This results in a typical deadlock situation.

15

4. The compiler which generated the KRs cannot recognize and cancel the interaction over time of the KRs so that no conflict situation arises.

20 The ratio between the cost of a circuit to be implemented and an optimum result should be as good as possible, i.e., the object of the invention is to provide a flexible, parallel, deadlock-free configuration that can be executed using moderate time and computing resources at a low cost. In this
25 context the following basic problems must be solved:

-if only KR3 were to be loaded, the process would be deadlock free but not optimum since KR5 could also be loaded.

30 -if KR3 is loaded but KR4 is not, and KR5 [is], KR4 must be pre-marked so that it has the highest priority in a subsequent loading sequence, which means high overhead.

Deadlock-free operation is ensured by the following procedure:

35

Improvements Through and Object of the Invention

The basic object of the present invention is a unit, hereinafter referred to as configuration table (CT), which has a hierarchical structure and may occur several times at each level, the number of CTs from the lowest hierarchical level to 5 the highest diminishing so that exactly one CT is present at the highest level. Each CT configures and controls independently from others and in parallel a plurality of configurable elements (CEls). CTs of the higher hierarchical levels can buffer configuration routines for lower-level CTs.

10 If more than one lower-level CT requires the same configuration routine, it is buffered in a higher-level CT and retrieved by the individual CTs, the higher-level CT retrieving the respective configuration routine only once from a global common configuration memory whereby a cache effect is 15 achieved. In addition to configurable modules, the present invention can be used as a cache procedure for instruction and data cache in microprocessors, DFP or the like having a plurality of arithmetic units. Some of the units to be described below may be omitted depending on the application 20 (e.g., FILMO) however, basically nothing is changed in the hierarchical structure. Therefore this application is considered a subset and is not described in detail. One considerable advantage of the method described over conventional cache procedures is that data and/or codes are 25 cached selectively, i.e., using methods adapted accurately to the algorithm.

The present invention also allows large cell structures to be reconfigured in a completely deadlock-free manner.

30

Description of the Invention

Instead of integrating, as previously, a central and global 35 unit in one module, with this unit processing all the configuration requests, there is a plurality of hierarchically (tree structure) arranged active units which can assume this

task.

A request from the lowest level (the leaves in the hierarchy) is forwarded to the next higher level only if the request
5 could not be processed. These steps are repeated for all the levels present until the highest level is reached.

The highest level is connected to an internal or external higher-level configuration memory which contains all the
10 configuration data required by this program run.

Due to the tree structure of the configuration units a kind of caching of the configuration data is achieved. Accesses to configurations mainly occur locally. In the most unfavorable
15 case, a configuration must be loaded from the higher-level configuration memory if the corresponding data is not available in any of the hierarchically arranged CTs.

Deadlocks are avoided in that a fixed time sequence of the
20 configurations to be loaded is introduced and the configurations are combined to form a list. The status information of the CELs is saved prior to loading and thus remains unchanged during the processing of the entire list of configurations.

25

Basic Principle of the CT

A configuration table (CT) is an active unit which responds to sync signals, known as triggers. The triggers are generated by
30 a two- or multi-dimensional matrix of electronic components usually for arithmetic or logical units, address generators, arithmetic units, and the like, hereinafter referred to as configurable elements (CEL). The trigger that occurs triggers a certain action within the CT. The task of the CT is to
35 assume the control of a plurality of CELs and to determine their arithmetic and/or logical operations. In particular, CELs must be configured and reconfigured. A CT assumes this

task by managing a plurality of possible configuration routines (KR) which in turn each contain a plurality of individual configuration words (KW) and configure [sic] a plurality of CELs on the basis of the trigger conditions using
5 one or more KR. Each CEL receives one or more configuration words, which are provided with the address of the CEL to be configured. A KR must be fully and correctly mapped to a plurality of CELs; a plurality of CELs can be combined to form groups; these groups are configured using different but fully
10 executed KRs. All CELs in a group are interconnected so that after a necessary reconfiguration has been determined it is communicated to all the grouped CELs, via a common signal (ReConfig), that each CEL must terminate data processing and go into a reconfigurable state.

15

Basic Principles of the Deadlock-free Reconfiguration

In systems that are reconfigurable in runtime the problem
20 arises that the system may enter a state in which two parts wait for each other, thus creating a deadlock situation.

This problem could be avoided by only loading a new configuration always entirely or not at all into the system or
25 by using a type of timeout procedure.

This entails a series of disadvantages (space requirements, running time, etc.) and problems such as:

- procedure if a configuration cannot be loaded.
- 30 - managing the order in which the configurations are loaded.
- deteriorated performance, since other configurations that could possibly be loaded into the CELs are not taken into account.

These problems can be eliminated with the method described in

the following. The starting point is a DFP system according to the related art.

A trigger signal is transmitted from a CEL to a CT. This CT
5 determines the trigger source and selects a configuration (KR) to be loaded using a lookup table. The incoming trigger signals are blocked, and no further triggers are accepted until the current configuration is completely processed. A configuration contains a plurality of commands which are
10 transmitted to a number of CELs. In a system configurable in runtime it is, however, not guaranteed that each configuration command (KW) can also be executed. Impossibility of execution may be caused, for example, by the fact that the addressed configurable element (CEL) has not yet completed its task and
15 thus cannot accept any new configuration data. In order to avoid performance deterioration, all configuration commands that could not be processed (since the respective CELs were not in a reconfigurable state and rejected the configuration (REJECT)) are written according to the FIFO principle in a
20 configuration command located in a special memory (FILMO) (to be described in detail below) after the last command. Subsequently the next configuration command is executed by the same procedure. This is repeated until the end of a configuration has been reached.

25 Subsequently, the CT assumes a state in which it accepts trigger signals allowing it to continue to load configurations. In this state, the CT processes the FILMO in regular intervals under the control of a timer.

30 The configurations to be loaded are prioritized by the fact that the CT runs through the FILMO memory before the configuration to be actually loaded is processed. A FIFO-like structure of the FILMO ensures that KWs which could not be
35 fully processed during previous trigger requests automatically obtain a higher priority with respect to the new [KW] to be processed. When the FILMO memory is processed each

configurable element (CEL) addressed by a configuration command is tested prior to or during the transmission of a KW to see whether it is in the "reconfigurable" state. If this state is "reconfigurable" (ACCEPT), the data is transmitted and deleted from the FILMO memory. If the state is "not reconfigurable" (REJECT), the data remains in FILMO and is processed again in the next run. The CT processes the next entry in FILMO.

This is repeated until the end of the FILMO is reached. Then the actual configuration activated by the occurrence of the trigger signal is processed. The structure of the FILMO corresponds to the FIFO principle, i.e., the oldest entry is processed first. In order to process the FILMO even when no new KR is loaded, the FILMO is run through in regular intervals under the control of a timer.

The other, uninvolved configurable elements (CEL) continue to work simultaneously during this phase without their function being affected. Thus it may occur that while the CT processes the FILMO, one or more configurable elements (CELS) change to the "reconfigurable" state. Since the CT with the processing may be at any location within the FILMO, the following may occur:

The CT attempts to process a first command whose addressed element to be configured (CEL) is not in the "reconfigurable" state. Therefore, the CT continues with the next command (KW). At the same time one or more configurable elements change to the "reconfigurable" state, among them also the configurable element which could have been written via the first configuration command. The CT processes a second configuration command (KW), which uses the same configurable element (CEL) as the first configuration command, however originates from another configuration. At this time, the configurable element (CEL) is in the "reconfigurable" state and the command can be processed successfully.

This no longer ensures that the configuration which should have been loaded first is actually completed first. Two partially completed configurations can now exist each of which requires configurable elements of the other configuration in order to be fully loaded. A deadlock situation has occurred, which is illustrated in Figure 18. Configuration A and configuration B are to be configured. The CT has already loaded the crosshatched part of configuration A and configuration B. Configuration A still requires the light double shaded area of configuration B for completion and configuration B still requires the dark double shaded area of configuration A for completion. Since neither configuration is fully completed and thus functional, the termination state in which one of the two configurations would be removed does not occur for either of the two configurations. Both configurations are waiting the for the release of the configurable elements still needed.

In the present method a deadlock is prevented in that the CT detects the states of all configurable elements prior to processing the FILMO and then allows no more changes or ignores the changes that occur until the sequence is completed. In other words, either the states of all configurable elements are saved prior to processing the FILMO or a status change is prevented while the FILMO is being processed. One possible technical implementation is the use of a register in each configurable element in which the status prior to processing the FILMO is saved. The CT only works on the basis of the states detected and not with the current states of the configurable elements. This ensures that each command (KW) to be processed is faced with the same state of the configurable elements (CELS). This step does not rule out that one or more configurable elements switch to the "reconfigurable" state while the FILMO is being processed. This change, however, is not immediately visible to the CT during processing, but only when the next run begins.

Configuration Sequences

In order to configure certain algorithms, it is absolutely necessary to strictly observe the order in which the KW are written in the CEL. For example, it is recommended that the bus system be configured first before a CEL is connected to a bus system, so that the CEL is not connected to a bus used by another routine. In other words, a CEL is only configured if the respective bus links have been previously configured.

10

In the method according to the present invention, a fixed sequence is observed as follows:

15 Configuration words (KWs) whose execution is decisive for the configuration of the subsequent KWS are especially marked (and referred to as KWR below). If the configuration of such a KWR fails, all subsequent KWS within the respective configuration routine (KR) are written to the FILMO and are not executed in this run. Even when running through the FILMO, all KWS located 20 in the sequence behind a KWR whose configuration failed are not executed in the current run.

Cache Procedure

25 The CT has a hierarchical structure, i.e., there are several CT levels in a module. The arrangement is preferably a tree structure (CT tree). An external configuration memory (ECR), containing all KR is assigned here to the root CT, while the configurable elements (CELS) which call up the individual KR 30 are assigned to the leaves. The configurable elements are assigned to the mid-level CTs of the same hierarchical level. A local internal memory is assigned to each CT. This memory is partially cleared when there is no more space for new KR to be stored or when this is explicitly requested by a special CT 35 command (REMOVE). The memory is cleared one KR at a time using a delete strategy such that in the ideal case only those KR are deleted that are no longer requested or are explicitly

identified in the REMOVE command. The number of KR_s deleted individually is such as to free the exact amount of memory space required to write the new KR_s to be loaded into the memory. As a result, as many KR_s as possible remain in the
5 memory to optimize the cache effect.

The advantage is that the each CT subordinate to any desired CT_x and thus located further up in the CT tree does not request a KR stored in the CT_x from the external configuration
10 memory ECR but obtains it directly from CT_x. This results in a cache structure over several levels. The data transmission overhead in the CT tree, in particular the required ECR memory bandwidth, is considerably reduced.

15 In other words, each CT buffers the KR_s of the CT below it. This means that the lower level CT_s obtain the required KR_s directly from those above them, and no memory access to the external ECR is necessary. Only if a required KR is not available in one of the higher level CT_s must the KR be loaded
20 by accessing the ECR. This results in a particularly efficient hierarchical cache structure for KR_s.

This structure also results in delete strategy options which however should be determined empirically, depending on the
25 application. These options include:

- deleting the oldest entries
- deleting the smallest entries
- deleting the largest entries
- deleting the entries retrieved least often.

35 Basic Principles of the Ct Hierarchies

In order to achieve a cache effect, CT_s are connected together

to form a hierarchical tree structure. A bus system (inter-CT bus) connecting an upper node (CT) to a plurality of lower nodes (CTs) is located between the individual nodes (CTs).
5 Lower nodes (CTs) request data from the upper nodes (CTs); the upper nodes then send data to the lower nodes. The lower nodes exchange status information using the networks between the higher nodes which therefore must resolve the addresses.

Ct Hierarchies and Addressing

10 CT hierarchies are arranged so that a binary tree can be used for addressing the individual CTs. This means that the least significant address bit identifies the individual leaves of the tree and each additional address bit selects the next
15 higher hierarchical level. Thus, each CT has a unique address.

The following table shows how the individual address bits are assigned to the respective levels:

20	3	2	1	0	Address Width
	-	-	-	*	Level 0: Leaves 1
	-	-	*	*	Intermed. Level 1 2
	-	*	*	*	Intermed. Level 2 3
25	*	*	*	*	Intermed. Level 3 4
					...

* = address bits used

- = address bits not used

30 If a higher level CT is assigned to a group of CTs, multiple address bits of the group are combined accordingly.

35 The following table shows how the individual address bits are assigned to the respective levels; a group with 8 CTs is

located on level 0. (Address bit 2-0) :

[Insert table]

5	4	3	2	...	0	Address Width
-	-	-	*	Level 0: Leaves	3	
-	-	*	*	Intermed. Level 1	4	
-	*	*	*	Intermed. Level 2	5	
*	*	*	*	Intermed. Level 3	6	
				...		

10

* = address bits used

- = address bits not used

15 The binary tree can have an unidimensional or multidimensional structure by having one binary tree for each dimension.

20 A certain CT (TARGET) is addressed in that the initiating CT (INITIATOR) either provides the exact target address or uses the relative TARGET address.

The evaluation of a relative address is described in more detail below.

25 Example of a relative address field for a two-dimensional addressing:

15	14	13	12	11	1	9	8	7	6	5	4	3	2	1	0
Higher CT	Broad-cast	± Y Address										± X-Address			

30

Bit 15 is set when the CT of the next higher hierarchical level is to be selected.

Bit 14 identifies broadcasts, i.e., selects all CTs.

The X/Y addresses provide the address of the TARGET based on the address of the INITIATOR.

5

The addresses are "signed" (plus or minus) integers. The TARGET is determined by adding the Y/X addresses of the address field to the current address position. Each level has a certain address width. The adders correspond to this width.

10

Overrun or underrun during addition means that the addressed CT is not below the current node and the address request is forwarded to the CT above it (the next higher node).

15

If no overrun or underrun occurs, TARGET is below the current node. The address bit calculated on the current level (see tables) selects the CT located directly below the current node. The next lower CT (node) is selected from this CT using the correspondingly calculated address bit.

20

Prioritizing Access in Ct Hierarchies

Access to the inter-CT bus is managed by an arbiter. All lower nodes have the same priority. The upper node has a higher priority. Therefore, an access transmitted from a higher node downward or one that has already covered a long distance from the INITIATOR is prioritized over other accesses.

Basic Structure of a CT

30

The following description of the CT provides an overview of the individual components. The detailed description of the main groups is given below.

35

The core of a CT is the control state machine (CTS) which controls all configuration routine (KR) processing. The garbage collector (GC) which controls the removal of KR from

the memory (CTR) of the CT; the FILMO which assumes the management of the KWs still to be processed, and the LOAD state machine which controls the loading of KR_s are assigned to the CTs.

5

The memory (CTR) is designed as a conventional read-write memory where all technically possible implementations may be used, and is used for local storage of KR_s for the respective CT and the CTs subordinate to it. As a special case, the
10 memory (CTR) can also be designed as a ROM, EPROM, EEPROM, flash ROM, or the like, in order to provide the module with a fixed, ASIC or PLD-type (see Background information) function.

10

In order to generate the CTR addresses, four pointers
15 configured as loadable counters are used:

1. Free Pointer (FP). Points to the first free memory location after the last KR in the CTR.

20

2. Garbage Pointer (GP). Points to an entry to be removed by the garbage collector (GC) from the CTR.

25

3. Move Pointer (MP). Points to the memory location in the CTR from which a valid configuration word (KW), not to be removed, i.e., an entry of a KR, is to be copied/moved to the entry defined by the GP.

4. Program Pointer (PP). Points to the KW currently being executed by the CTs.

30

KWs are forwarded to the respective CELs via an output interface (OUT). The CELs acknowledge (ACCEPT) the receipt of the KW, provided they are in a reconfigurable state. If a KW is not acknowledged (REJECT) it is temporarily buffered in a
35 FIFO-like memory (FILMO) to be written again to the addressed CEL at a later time without using the program pointer.

The CTS receives, via trigger signals, a request to process a KR. The trigger signals pass through a mask, i.e. a filter, which filters out (masks out) undesirable triggers. A mask can be designed according to the related art using an AND gate,
5 which links a trigger with an AND enable signal. The triggers are converted into binary signals via a prioritized round-robin arbiter (SCRR-ARB). A prioritized round-robin arbiter combines the advantage of the democratic character of a round-robin arbiter with the recognition of the next enable in a
10 cycle, i.e., the advantage of a priority arbiter.

The masked triggers are switched as an address to a first lookup table (LUT1), i.e., a memory that assigns the ID of the respective KR to the trigger incoming as an address and
15 outputs it over the data lines.

In a second lookup table (LUT2) the ID of the KR is assigned to the address of the memory location of the KR in the CTR. The second lookup table is used not only for assigning trigger
20 signals, but commands using an ID as a parameter also use the LUT2 for address assignment.

The assignment of the trigger signals to the respective IDs is entered into LUT2 via the "REFERENCE" command to be described
25 later. LUT2 is managed, i.e., the IDs are assigned to the addresses in CTR, automatically by the CTS and the GC.

To better understand the CT, a possible basic set of commands is illustrated below:

30

1. BEGIN <ID>

BEGIN <ID> denotes the start of a configuration routine. <ID> provides the unique identification number of the configuration routine.

35

2. STOP

STOP denotes the end of a configuration routine. At this point

the configuration table (CT) terminates the processing of the configuration routine. The garbage collector (GC) terminates the removal of entries of this configuration routine.

5 3. EXECUTE <ID>

Jumps to the start (BEGIN <ID>) of a configuration routine. If this routine is not present in the CT memory, it is requested from the next higher CT, or loaded from the memory.

10 4. LOAD <ID>

Requests the KR <ID> from the next higher CT.

5. REMOVE <ID>

Calls the GC in order to remove the configuration routine <ID>
15 from BEGIN <ID> to STOP from the CT memory and pre-write the subsequent configuration routines to the point where no memory hole is created by the removed configuration routine.

6. PUSH <FORCED> <ADDRESS> <DATA> <EXIT>

20 Writes the configuration data <DATA> to the <ADDRESS> register. If <FORCED> is set, data is written even if the RECONFIG flag of the respective target register is not set. <EXIT> is used to indicate a KWR which interrupts further execution of the subsequent KWRs in the event of a REJECT.

25

7. MASK <SR> <TRIGGER>

Sets or resets the trigger mask using <TRIGGER> as a function of <SR> (Set/Reset).

30 8. WAIT <UNMASKED> <TRIGGER>

Stops the processing of the configuration routine and waits for <TRIGGER>. If <UNMASKED> is set, there is a response to the expected trigger regardless of the state of the trigger mask.

35

9. TRIGGER <TRIGGER> <CT#>

Sends the binary value of a trigger to the higher level CT

addressed by CT#.

10. GETBUS/GETCTS

Establishes a connection to the inter-CT bus.

5

11. LOOSEBUS/LOOSECTS

Terminates the connection to the inter-CT bus.

12. REFERENCE <TRIGGER> <ID>

10 Writes the value <ID> into the LUT2 at the address <TRIGGER>, whereby a certain configuration routine (KR) is assigned to a trigger signal.

15 The commands EXECUTE, LOAD, REMOVE, PUSH, MASK, WAIT, TRIGGER, REFERENCE are only valid within the BEGIN - STOP brackets. The commands are not executed outside these brackets.

A configuration routine (KR) has the following structure:

20 BEGIN <ID>;

...

valid commands

...

STOP;

25

Indirect Addressing (Referencing)

30 The cache principle of the CT allows a KR to be buffered in a CT, the KR being used by a plurality of different lower-level CTs or CELs.

35 If the lower-level units access the external interface of the module (e.g., RAM, periphery), it becomes necessary to store different addresses or parts of the external interface. This would make the contents of the required individual KRs fundamentally different. Caching is no longer possible.

Indirect referencing provides a remedy. For this purpose, special KR (hereinafter referred to as IKR) are used, which contain and set the required external parameters. Other different KRs may be called up via triggers on different hierarchical levels. The actual KR is called up starting at the end of an IKR. The IKR are not cachable, however, while the KR called up are all uniform and are therefore cachable. It is advisable that the size of the IKR be reduced to the absolute minimum, i.e., only the external and different parameters and the call up of the uniform KRs.

An indirect configuration routine (IKR) is structured as follows:

```
15 BEGIN <ID>;
...
xxx; valid commands where only external peripherals are to be activated,
TRIGGER <ID>; start, stop, or load requests to peripheral
20 processes
...
GOTO <ID>      Jump to uniform KRs
STOP;
```

25

Special Cases:

1. WAIT_FOR_BOOT

This command is only valid at the first address of the CTR. During the boot sequence, the complete boot KR is initially written into the CTR, but not the start sequence of the boot KR BEGIN <0>. In its place (at address 1) is WAIT_FOR_BOOT which is set automatically after a RESET. WAIT_FOR_BOOT is not overwritten with BEGIN <0> and the CTS does not begin processing the boot KR until the entire boot KR has been written to the CTR.

WAIT_FOR_BOOT should not occur within a program.

2. BOOT <CT ID>

BOOT <CT ID> marks the CT in which the following boot KR should be written. No BEGIN follows after BOOT <CT ID>, the boot KR is not terminated by STOP, but by a subsequent BOOT <CT ID>. A STOP ends the boot sequence.

5 BOOT <CT ID> should not occur within a program.

Boot Sequence

10 After a RESET, the CT of the uppermost hierarchical level (ROOT CT) loads the boot KR into the CTs of the lower hierarchies. For this purpose, there is a jump to a fixed address (BOOT ADR) in the external configuration memory (ECR) 15 assigned to the ROOT CT. The ROOT CT executes this jump and arrives at the boot sequence. This is structured as follows:

BOOT <CT ID0>; COMMAND; COMMAND; ...
BOOT <CT ID1>; COMMAND; COMMAND; ...
20 BOOT <CT IDn>; COMMAND; COMMAND ...
STOP;

25 During the boot sequence the complete boot KR is initially written into the CTR starting at address 2 of the CT defined by <CT ID>. The start sequence of the boot KR (BEGIN <0>) is not written to address 1. In its place is WAIT_FOR_BOOT which is automatically set after a RESET. STOP is not written to the end of the boot KR in the CTR, overwriting WAIT_FOR_BOOT with BEGIN <0> until the entire boot KR has been written into the 30 CTR and the ROOT CT has reached the next BOOT <CT ID>. The CTS starts processing the boot KR.

Loading a Configuration Routine

35 There are three basic mechanisms for requesting a configuration routine in addition to the boot KR:

1. Execution of a LOAD <ID> by the CTS

2. Execution of an EXECUTE <ID> by the CTS, the KR with the respective ID not being available in the CTR.

5

3. Occurrence of a trigger which is translated into an <ID> whose respective KR is not available in CTR, via LUT1.

The sequence is the same in all three cases:

10

The ID of the requested KR is provided as an address to LUT2. LUT2 checks whether a valid address exists in CTR. If it does not exist, i.e., <ID> points to the value 0 in LUT2, load <ID> is sent to the CTS.

15

Subsequently CTS requests the <ID> of the respective KR from the hierarchically higher-level CT. This request reaches the higher-order CT in the form of a trigger and is evaluated by it accordingly.

20

The higher-order CT transmits the requested KR to the requesting CT. The data is written, starting at the address to which FREE-POINTER (FP) points, into the CTR, the FP being incremented by one after each write access.

25

When FP reaches the upper limit of the CTR, the garbage collector (GC) is called up to remove the lowermost KR within the CTR and to compress the CTR. The FP is reset. This sequence continues until the KR to be loaded fits completely

30

into the CTR.

Jump Table in the Configuration Memory

The configuration memory assigned to ROOT CT contains all the 35 KR that must be loaded for an application. In the external configuration memory (ECR) there is a jump to the boot configuration routine at a fixed address (ADR BOOT). At

another fixed address range (LUT ECR) of any desired length which is, however, predefined within an application there are jumps to the individual KRs. The <ID> of each KR is used as an address in ECR where the start address of the respective KR is located; thus KRs are addressed indirectly:

ID -> LUT ECR -> KR

Modifying the Kr in the Configuration Memory

The KR with the ID <A> is to be modified. The HOST initially writes the new KR for the ID <A> to a free memory location in the ECR. The ID <A>, together with the new address of the KR, is written in the configuration memory from the higher-level unit (HOST) into a ROOT CT register provided for that purpose. Subsequently, when a STOP is reached or during IDLE cycles, i.e., as long as no KR is being executed, all CTs remove the KRs related to this ID from the CTR and set LUT2 at address <A> to "NoAdr," which means there are no valid address entries for ID <A> in LUT2. If the ID <A> is requested again, the missing entry ("NoAdr") at the location <A> in LUT2 forces each CT to request the KR <A> from the ECR again.

25 The FILMO

A KR is composed basically of the command PUSH, which writes new configuration words to a certain address. If a configuration word of the type KW cannot be written because the configurable element (CEL) addressed is not ready to receive a new configuration (REJECT), the configuration word is written into a memory hereinafter referred to as FILMO instead of to the configurable element (CEL) addressed. The subsequent commands are processed normally until a configuration word cannot be written again; it is then written into the FILMO.

If a configuration word of the type KWR cannot be written because the configurable element (CEL) addressed is not ready to receive a new configuration (REJECT), the configuration word is written into a memory hereinafter referred to as FILMO instead of to the configurable element (CEL) addressed. All subsequent commands to the end of the KR are written directly into the FILMO, rather than to the CEL.

The entire FILMO is run through in IDLE cycles and before each execution of a new KR. Each word of the FILMO that is read out, beginning with the oldest data word is to be sent to its addressed element according to the known FIFO principle; the addressed element must be ready to receive the configuration word. If the data words can be written from the beginning (i.e., the addressed configurable elements (CELS) are ready), the entry is removed from the FILMO by the FIFO principle. If a configuration word cannot be written, it is skipped and is not removed from the FILMO. Contrary to FIFO, the data after the skipped configuration word continue to be read. The configuration words written after a skipped configuration word can, depending on the implementation of the FILMO, be

1. marked as written and not deleted from FILMO; configuration words marked as written are no longer read during subsequent runs or are immediately deleted if a skipped configuration word is no longer located in front of them;
- or
2. deleted from the FILMO; the configuration words before and after the deleted configuration word are preserved. The subsequent words must be shifted forward (up) or the previous words must be shifted backward (down) for deletion, the order of the configuration words being preserved in all cases.

If a new KR is executed, the configuration words (KW) which could not be written by the CTS to the addressed elements (CELS) are appended again to FILMO, i.e. the KW are written to the end (in the reading direction) of the FILMO. If the FILMO is full, i.e., there are no free entries for configuration

words, the execution of the KR is stopped. The FILMO is run through until a sufficient number of configuration words could be written and a corresponding number of free entries have been created, whereupon the KR continues to be processed.

5

The FILMO is a FIFO-like memory which is always run through linearly starting from the oldest entry; contrary to a FIFO, however, entries are skipped (First In Linear Multiple Out).

10 Function of the Configuration Table State Machine (Cts)

The configuration table state machine (CTS) is responsible for controlling the CT. It executes the commands of the KR and responds to incoming triggers. It manages the FILMO by reading
15 the FILMO during IDLE cycles and before executing a KR.

It responds to the signals illegal <TRG> (Illegal Trigger, see Figure 1, 0102) and load <ID> generated by the LUT structure.
load <ID> is generated when a cache miss exists in LUT2 (0105)
20 or the KR/IKR referenced by ID has been marked as deleted (0107). It responds to the control signals of the higher level CT.

An example of implementation for processing the commands is
25 illustrated in Figures 2 through 7.

Control Signals to Higher Level Cts

- 30 - illegal <TRG> (0102)
 shows to the higher level CT that an unknown trigger <TRG> has occurred.
- load <ID> (0105/0107)
 requests the higher level CT to load the <ID>.
35 - trigger <TRG> <CT#> (0108)
 sends a trigger <TRG> to the higher level
 or

to the addressed CT <CT#>.

Control Signals from Higher Level CTs

- 5 - remove <ID> (see Figure 15, 1513)
 requests the CT to delete the <ID>.
 - write_to_FP <data> (see Figure 2, 0205)
 sends the data to the CT. The data is appended to the end of
 the data in the memory.

10

Function of the Garbage Collector (Gc)

The CTR has two problems:

- 15 1. If a LOAD or EXECUTE command or a trigger refers to an ID
 whose KR is not present in the CTR, the KR must be re-loaded.
 Sometimes however there is not enough space available in the
 CTR to load the requested KR.
- 20 2. If a REMOVE <ID> occurs, the respective KR must be removed
 from the CTR. This creates a gap unless the KR is located at
 the end of the CTR. When a new KR is loaded, the gap may not
 be completely filled again or the gap may be too small for the
 new KR. This results in fragmenting of the CTR. The function
25 of the garbage collector is to remove KR from the CTR in order
 to create space for new entries AND to reorganize the CTR
 after the removal of entries so that all remaining KR are
 located in succession in the memory as a contiguous block and
 the freed memory blocks are located at the end of the CTR as a
30 contiguous block.

Thus new KR can be re-loaded in an optimum manner and without
loss of memory space.

35 Evaluating Trigger Pulses

Each CT has a link to a plurality of trigger signals belonging

to its respective hierarchical level, which are combined to form a bus. Incoming triggers are evaluated via a mask, i.e. only the enabled trigger signals are forwarded. The enabled trigger signals are buffered in a sample register

5 synchronously with the system clock (sampled). An arbiter selects one of the stored trigger signals and converts the signal into a binary vector. The selected trigger signal is deleted from the sample register. The binary vector is forwarded to a first lookup table (LUT1) which translates the

10 binary vector into the identification number (ID) of the configuration routine (KR) to be called up. The ID is translated into the address of the KR in the CT memory (CTR) in a second lookup table (LUT2). The CT state machine (CTS) sets its program pointer (PP) to this address and begins the

15 execution of the KR. The prerequisite is that each trigger enabled via the mask has a corresponding entry in LUT1. Otherwise an error condition is forwarded to the CTS (illegal trigger); each ID = "NoAdr" is evaluated as a non-existing entry. "NoAdr" is a token selected depending on the

20 implementation.

If there is no entry in LUT2, i.e., the KR of the respective ID is not in the CTR, a load request is sent to the CTS (load <ID>).

25

Sending Trigger Pulses to the Higher Level CT

In addition to the interface described above to a higher level CT for loading KR, there is an additional interface for exchanging freely definable commands, in particular, however, trigger vectors. A CT sends

- either a command to all the other CTs (BROADCAST)
 - or a command to any addressed CT (ADDRESSED).

35 The "Triggervector" command represents a binary value that refers to an entry in the LUT2 of the receiving CT.

It is necessary to send trigger vectors in order to start a KR within an IKR in an additional CT, in order to activate the periphery or the memory, for example.

- 5 In order to forward trigger vectors to a higher level CT,
there are two mechanisms:
- 10 1. A bit indicating whether the content of the memory is
regarded as a KR ID or as a binary value for a trigger pulse
is supplied to LUT1. If a trigger pulse is present, the data
content of LUT1 is sent directly to the higher level CT as a
trigger.
- 15 2. The binary value of a trigger can be indicated with the
TRIGGER command which is sent directly to the higher level CT.
(As an alternative, IDs can also be transmitted directly
instead of a trigger value).

20 In order to start a KR in another CT via trigger vectors, a
synchronization procedure must be created in order to prevent
a deadlock. The procedure must make sure that only one KR
within a certain group of CTs starts other KRs in other CTs
within this group. Starting several KRs at the same time may
result in a deadlock among the CTs, similar to the deadlock on
25 the CEL level described above.

In principle such a procedure runs as follows:

A KR is structured as follows:
30 ...
GETCTS/GETBUS
TRIGGER <ID>, <CT#>
TRIGGER <ID>, <CT#>
...
35 LOOSECTS/LOOSEBUS
...

The "GETCTS" command within a KR of the CT (INITIATOR) indicates that signals are sent to other CTs (TARGET) in the following. With Trigger <ID>, <CT#> the ID of a KR to be started is sent to the CT having the unique ID CT#. The
5 trigger is initially sent to the directly higher level CT, which sends the trigger according to the CT# to a lower level CT within its CT space or to its higher level CT (see CT hierarchies). When the command reaches the TARGET, the latter acknowledges receipt.

10

When the command runs through a CT, a priority ID of the command is incremented by one. If the forward request of a command encounters another request within a CT, the command having the lowest priority is rejected. Thus

- 15 a) it is ensured that only one command propagates at any one time within an overlapping system and thus only one KR is started, which prevents deadlock as required;
b) it is ensured that the command that has propagated least far is rejected, which may result in enhanced performance.

20

After a command has been rejected, all previous commands within the GETCTS/LOOSECTS are also rejected, i.e., INITIATOR sends the DISMISS signal to all TARGETs and the execution of the KR starts again after a wait time upon GETCTS.

25

The acknowledgments of all triggers within a GETCTS - LOOSECTS command segment are sent to the INITIATOR CT. Processing of the next command continues with each incoming acknowledgment.

30

When the LOOSECTS command is reached, INITIATOR sends the GO signal to all TARGETs. This starts the execution of the KR having the ID transmitted by the trigger by the TARGET CTs.

35 After the occurrence of a trigger, TARGETs change to a state in which they wait for the occurrence of a GO or DISMISS signal.

For better implementability, a slightly modified procedure is also presented:

5 A bus system (inter-CT bus) is located between the CTs in a group of a hierarchical level. This bus system connects all CTs of the group and a CT of the next higher level group.

10 The bus system is arbitrated by a CT through the GETBUS command which is functionally similar to GETCTS. The commands are forwarded via the bus system to the CTs of the same group. If the addressed CT# is not within the group, the higher level bus is automatically arbitrated by the higher level CT and the command is forwarded. The arbitrated buses remain assigned to INITIATOR and thus blocked for all other CTs until either a
15 reject occurs or the LOOSEBUS command frees the bus. LOOSEBUS is comparable to LOOSECTS. Before executing the LOOSEBUS command, the GO signal is sent to all involved CTs. This is accomplished either through the LOOSEBUS command or through a special command that precedes it. Commands, in particular
20 triggers, are also processed according to the basic procedure described above. A reject occurs if a bus system cannot be arbitrated. The CTs of a certain level have the same priority for arbitration; the higher level CT has a higher priority.
When a command is sent via the inter-CT bus, the command
25 remains active until the addressed CT accepts (ACCEPT) or rejects (REJECT) the command.

The Prioritized Round-robin Arbiter

30 The prioritized round-robin arbiter (single-cycle round-robin arbiter SCRR-ARB) has a clock synchronous structure, i.e. for each (positive or negative depending on the implementation) cycle flank (TF1) it delivers a result. The incoming signals (ARB-IN) pass through a mask (ARB-MASK) which is managed by
35 the arbiter itself by the procedure described below. The output signals of the mask are supplied to a priority arbiter (ARB-PRI0) according to the related art. The arbiter delivers

a result (ARB-OUT), synchronized with the system clock, for each cycle flank (TF1), i.e., the binary value of the highest-priority signal after the mask (ARB-MASK). A signal (VALID) which indicates whether the binary value is valid or invalid

5 is assigned to the result. Depending on the implementation of the priority arbiter, it is possible that the same binary value is generated when the signal 0 is present and when no signal is present; in this case VALID indicates that the result is invalid if no signal is present. This signal is

10 1. output as the result of the arbiter
and

2. sent to a decoder which decodes the binary values as shown in the following table for a three-bit binary value as an example. (The coding procedure can be adapted to any desired

15 binary value according to this principle):

Binary Value (ARB OUT)	Decoding (ARB DEC)	Note
111	0111 1111	
110	0011 1111	
101	0001 1111	
100	0000 1111	
011	0000 0111	
010	0000 0011	
001	0000 0001	
000	1111 1111	Reset state and when binary value (ARB OUT) invalid

30 A register (ARB REG) which receives the decoded values (ARB DEC) of the decoder at a cycle flank (TF2) that is the inverse of TF1 is assigned to the decoder. ARB DEC is fed back to the mask (ARB MASK) and enables the individual input signals (ARB IN).

The functional sequence in the arbiter is as follows:

1. After a RESET, all ARB IN are enabled via ARB MASK, since ARB DEC sets all signals to "enable".

5

2. The highest priority ARB IN set (for example signal 7 (binary 111) has the highest priority and 0 (binary 000) the lowest priority in the above table) is output as a binary value.

10

3. Via ARB DEC the signal is blocked, as are all the other inputs that may have had a higher priority, but are not set.

15 4. The following steps 5 and 6 are repeated until signal 0 (binary 000) is reached or no signal is present after ARB MASK. Then ARB DEC (see decoding table) enables all signals through ARB MASK via ARB DEC again and the sequence starts at step 2.

20 5. The ARB IN set that now has the highest priority is output as a binary value.

6. Via ARB DEC the signal is blocked, as are all the other inputs that may have had a higher priority but are not set.

25 (Continues with step 4)

Thus it is achieved that all input signals are handled with the same priority and one of the input signals (ARB IN) is decoded as a binary signal and output (ARB OUT) in each clock cycle. ARB REG can be provided with an enable input (EN) which allows the contents of the register to be changed only at TF2, when a corresponding signal is present. The result is that a binary vector is not output in each cycle, but is output instead as a function of an enable by EN and TF2. The input is needed for synchronization when the downstream circuit cannot perform the processing in one clock cycle, but needs multiple cycles before accepting the next binary vector.

Sometimes it is recommended that the arbiter consider a series of signals as having a higher priority, while the majority of signals have the same priority. This is necessary, for example, in the case of the previously described procedure for forwarding signals between CTs. In order to give a signal a higher priority, the link of the ARB PRIO having the highest priority is not masked, i.e., it bypasses the mask (ARB MASK). Thus the signal is treated preferentially.

10

Structure of a Microcontroller-based CT

Contrary to the previous descriptions, a CT can also be implemented to have a microcontroller architecture.

It can be easily seen that the basic functions such as trigger control, lookup tables LUT1 and LUT2, as well as the inter-CT communication and writing the KW to the CEL can be readily accomplished by a microcontroller. Only the structure of an efficient FILMO represents a problem which is mainly manifested in the performance that can be achieved. Therefore the structure of the FILMO will be described separately.

25 Structure of the FILMO

The FILMO is not designed as a separate memory. Instead, the conventional program memory is extended to include FILMO functionality. For this purpose, an additional bit (FILMO BIT) which indicates whether or not the corresponding KW has been written into the CEL is assigned to each KW. If the FILMO BIT is set, the corresponding KW is not executed. When a KW is written into the memory, the FILMO BIT is reset. All KRs within a CT are connected via a chaining list (FILMO LIST) in the order in which they were called up by triggers or LOAD <ID>. A KR remains in the FILMO LIST until it is fully executed, after which it is removed from the list. The FILMO

LIST is run through according to the FILMO procedure and thus represents a direct substitute for the FILMO memory.

5 (For the sake of completeness we shall note that contrary to the original FILMO procedure no KR may occur twice in the list. If a KR that is still in the FILMO LIST is called up, its execution must be delayed until it is removed from the FILMO LIST.)

10 The structure of a FILMO memory location is as follows:

FILMO BIT	KW
-----------	----

15 Commands

The microcontroller supports the following commands that have a direct influence on the FILMO:

20 PUSH write a KW to a CEL

PUSHSF write a KW to a CEL and set the FILMO BIT if the KW has been accepted (ACCEPT)

25 PUSHRET write a KW to a CEL and return (RETURN) from the subroutine if the KW has not been accepted by the CEL (REJECT). This command is used when subsequent KW in the KR are dependent on the configuration of this KW (ACCEPT); their configuration is prevented by the return from the KR until PUSHRET is successful (ACCEPT).

30
35 PUSHNR write a KW to a CEL only if no REJECT occurred previously within the KR. Like PUSHRET, PUSHNR is also used to handle dependencies in the order in which the KWs are configurated.

Garbage Collector

As described previously, a garbage collector (GC) is used to remove KR's that are no longer needed. The GC is started either when there is insufficient space in the memory for loading a new KR and IDs must be removed; or when a KR is explicitly deleted by the REMOVE command with the indication of the ID of the KR to be deleted.

- 10 In order to configure the GC run as simply as possible, all KR's are connected via a chained list. GC runs through the list and removes the KR's that are no longer needed by overwriting them with other KR's and adjusting the list entries accordingly. All the remaining KR in the memory are shifted so 15 that the memory gap created by the KR to be deleted is closed and a larger contiguous free space is created at the end of the memory.

Structure of a KR

20 One possible basic structure of a KR is shown in the following table:

25	jmp START;
	length
	garbage - previous
	garbage - next
30	FILMO - previous
	FILMO - NEXT
	CACHE - statistics
	KR - statistic
	START;
	...
35	ret;

At the beginning of the KR, a jump occurs over the following

header to the start of the command sequences. It follows the double-chained list for the garbage collector in which all KR are linked to one another. "length" indicates the length of the KR. This information can be used for Block-Move commands according to the related art which are used when the KR have to be moved in the memory (Garbage, Load, etc.).

The FILMO is formed in the following double-chained list; only those KRs that contain KWs which have not yet been written to the CEL are connected to one another.

Statistics on the condition of the cache follow, containing, for example, the number of KR calls (the value is incremented by one for each call), the age (can be measured using the number of GC runs over the KR; the value is incremented by one for each GC run), etc. The GC can evaluate these statistics when a KR must be removed for reasons of memory space.

Considerable advantages result for the caches from these statistics. Thus, for example, the microcontroller can be programmed depending on the cache algorithm used according to the requirements of the application so that

1. the oldest/newest KR
2. the smallest/largest KR (see the "length" entry)
3. the least often/most often called KR

is deleted from the cache when free memory is needed. Of course, additional appropriate status information can be stored. Such a selective caching is not possible with the cache structures known today. In particular, freely programmable cache algorithms are not supported in caches according to the related art.

In conclusion, there are KR statistics containing, for example, the number of not yet configured (REJECT) or configured (ACCEPT) KWs. At the same time, the address of the first KW yet to be configured can be stored. This has the advantage that in the case of a FILMO run a jump can be made directly to the KW and the complete KR does not have to be run

through, which results in a considerably enhanced performance.

In conclusion, it should be noted regarding the KR that the chained lists are preferably built up by entering the previous/next ID, since thus the absolute memory addresses can be easily shifted by the GC. Only relative jumps instead of absolute jumps should be used within a KR in order to avoid problems when loading the KR and during GC runs, since the absolute addresses are modified in those cases.

10

For the sake of completeness it should also be mentioned that, according to the above-described principle, the FILMO is run through prior to executing a new KR (upon a trigger or command even from another CT) and the status of the CEL (reconfigurable or not) is saved prior to running through the FILMO, also when a microcontroller is used.

15

Figures

20 The figures described below show the management of configuration data according to the method presented, with reference to an example of implementation:

25 Figure 1: address generation procedure within the lookup tables

Figures 2-7: processing of the commands and function of the state machines

30 Figure 8: Structure of the SCRR-ARB

Figure 9: Structure of the LUT1 and LUT2

35 Figure 10: Structure of the pointer arithmetics and the CTR

Figure 11: Structure of a FILMO

Figure 12a: Hierarchical arrangement of the CTs

Figure 12b: Sending a trigger between the CTs

5 Figures 12c, d: Method for sending a [incomplete]

Figure 13: Calling up a KR by multiple IKR

10 Figure 14: Structure of the LUT1 of a ROOT CT

Figure 15: Structure of the HOST control of a ROOT CT

Figure 16: Illustration of the LUT and ECR principle

15 Figure 17: Sequence control of a CT of mid-level hierarchy and of a ROOT-CT

Figure 18: Deadlock problem during the configuration of a two-dimensional array (see patent description)

20 Figure 19: Illustration of the FILMO principle

Figure 20: Basic principle of the inter-CT communication

25 Figure 21: Example of implementation of the inter-CT communication according to the GETCTS method

Figure 22: Example of implementation of the inter-CT communication according to the GETBUS method

30 Figure 23: Bus structure of the inter-CT bus

Figure 24: Addressing within CT hierarchies

35 Figure 25: GARBAGE list

Figure 26: FILMO list

Figure 27: FILMO function within a KR

Figure 28: Storing the states prior to the execution of a KR or the FILMO.

5

Description of the Figures

Figure 1 shows the sequence of the CTR address generation within a CT. An incoming binary trigger vector (0101) is
10 translated into a valid KR or IKR ID in LUT1. If no valid ID exists, an "Illegal Trigger" signal is generated (0102), which indicates that the trigger is not known in LUT1. The signal can be forwarded to the higher level CT as an error message or
15 ignored. The translation of "triggers" according to "ID" is entered into LUT1 using the "REFERENCE" command.

A valid ID (0103) is forwarded to LUT2. IDs provided within commands, i.e. through an operand (0104), are sent directly to LUT2. LUT2 translates an incoming ID to the address of the
20 KR/IKR within the CTR. If the KR/IKR is not stored in CTR (there is no cache) the "miss" signal is generated (0105). If the translated address of the KR/IKR is marked with the "NoAdr" token, it is indicated with "NoEntry" (0107) that the address has been deleted. "Miss" and "NoEntry" indicate that
25 translation into an address internal to CTR is not possible. On the basis of this signal the LOAD state machine loads the KR/IKR having the corresponding ID from a higher level CT.

30 If a valid address is available, it is forwarded to the pointer arithmetics of the address generator (0106). An incoming binary trigger vector is translated in LUT1 either into an ID or into another trigger vector; in this case the trigger vector is output (0108).

35 Figure 2 shows the sequence when loading a KR/IKR. The ID (0201) of the KR/IKR to be loaded is initially sent to the higher level CT. Then the value of FreePointer (FP) is entered

in LUT2 at the location of the entry for the requested ID. FP points to the entry after the last entry used for a KR/IKR in the CTR. This is the first entry, on which the KR/IKR to be loaded is stored.

5

The state machine waits for a data word from the higher level CT. As soon as the word is available, it is written to the location referenced by FP. FP is incremented. If FP points to an entry after the end of the CTR, the first entry in the CTR is removed in order to create room (0202); FP is updated. If the data word sent by the higher level CT is "STOP," the loading procedure is interrupted (0203); otherwise the state machine continues to wait for a new data word (0204).

10 15 Figure 3a shows the "MASK" command. The operand of the command is written into the MASK register. The MASK register is located at the input of the trigger signals before LUT1 and masks out invalid triggers.

20 25 In Figure 3b, the operand of the command is sent to the other CTs as a trigger vector by the "TRIGGER" command.

In Figure 3c, the translation of a trigger into the corresponding KR/IKR ID is written into LUT1 using the "REFERENCE" command.

30 35 Figure 4a shows the "WAIT" command. The operand of the command is written into the WAITMASK register. All triggers except for the one/ones expected and therefore enabled in WAITSMASK are ignored. Only upon the occurrence of the trigger does the system return to the program flow.

Figure 4b shows the "PUSH" command. The configuration word is sent to the addressed configurable element (CEL). If the CEL does not accept the configuration word, for example, because the CEL is in the "non-configurable" state, the configuration word is written into the FILMO (0401).

Figure 5 shows the sequence of a "REMOVE" command. There are two call variants:

1. The first KR/IKR located in the CTR is removed from the CTR. Address 0 of the CTR is assigned to the GarbagePointer (GP).
5
2. A specific KR/IKR defined by its ID is removed from the CTR. The first address of the KR/IKR to be removed from the CTR is assigned to the GarbagePointer (GP) (0502).
10

The MovePointer is loaded with the value from the GP. GP and MP refers to a "BEGIN <ID>" command in the CTR even if the first KR/IKR is to be removed from the CTR. The relevant ID is marked as invalid in LUT2. MP is incremented until the "BEGIN <ID>" of the next KR/IKR located in the memory is reached (0503), OR MP is equal to FreePointer (FP), which means that the KR/IKR to be removed is the last one in the CTR (0504).

20 - In this case, FP is loaded with the value from GP, whereby the memory locations occupied by the KR/IKR to be deleted are marked as free; the "REMOVE" function is terminated (0505).

25 - Otherwise ("BEGIN <ID>" is reached (0506)) the data referenced by MP is copied to the memory location referenced by GP. MP and GP are incremented. This procedure continues until MP reaches the end of CTR or the position of FP (0507). If during the sequence a memory location containing "BEGIN <ID>" is referenced by MP, the entry for the respective ID is 30 overwritten in LUT2 with MP (0508); thus the correct memory location is output in the case of a lookup.

Figure 6 shows the sequence diagram of the FILMO. A FILMO contains three pointers:

- 35 1. WriteP: the write pointer of the FILMO RAM;

2. ReadP: the read pointer of the FILMO RAM;
3. FullP: the status pointer representing the "filling status" of the FILMO RAM that prevents overrun or underrun.

5

A one-bit register "BeginF" indicates whether the current read access is at the beginning of the FILMO RAM (TRUE), i.e. there are no deleted entries between the read pointer and the beginning of the FILMO RAM; or the read pointer is in the middle of the FILMO RAM (FALSE), i.e. there are entries used between the read pointer and the beginning of the FILMO RAM. Furthermore, there are two registers for storing the states of ReadP and FullP. It is necessary to save both registers upon the occurrence of the first undeleted entry, because a subsequent read access to the location of this entry must be begun with the readout. On the other hand, ReadP and FullP must be modified during the current read sequence in order to obtain the next read addresses, or to establish the end of the FILMO RAM. The beginning and the end of the memory cannot be established using a zero address or a maximum address due to the structure of the FILMO as a FIFO-like structure ("ring memory"). Two sequence paths originate from the base state:

1. Read path (0601)
25 FullP and ReadP are saved in the registers. The processing loop begins:

BeginF is TRUE.
30 If FullP is equal to 0, ReadP and FullP are read back from their registers (0602) and the state machine jumps back to its base state.

Otherwise (0603) it is tested to determine whether the entry in the FILMO to which ReadP is pointing is equal to "NOP," i.e., whether this is an entry marked as deleted in the middle of the FILMO. If this is not the case (0604) it is attempted

to write the entry into the configurable element (CEL). If this is not successful (REJECT, 0605), because CEL is not reconfigurable, BeginF is set to FALSE, FullP is decremented, and ReadP is incremented. The state machine jumps to the beginning of the processing loop (0606).

If the write of the entry to CEL (0607) is successful or the entry is an NOP, BeginF is tested: BeginF == TRUE (0608): there are no undeleted entries before this one. FullP is incremented, ReadP is saved in the associated register in order to secure the new beginning of the FILMO. FullP is saved in order to secure the current data set; ReadP is incremented.

BeginF == FALSE (0609): FullP is incremented and the current entry in the FILMO RAM is overwritten with NOP, i.e. the entry is deleted. ReadP is incremented.

In both cases the state machine jumps to the beginning of the processing loop.

20

2. Write path (0610)

It is tested whether the FILMO RAM is full by checking FullP for the maximum value. If this is the case (0611), a jump is made to the read path in order to create room.

Otherwise the data word is written into the FILMO RAM and WriteP and FullP are incremented.

30 Figure 7 shows the sequence in the main state machine. The base state (IDLE) is exited as soon as a

35 1. REMOVE command from the higher level CT occurs (0701): the REMOVE command is executed and the state machine returns to IDLE.

2. A trigger signal for generating a trigger occurs between

the CTs (0702) :

The trigger is output.

5 The state machine jumps to the "STOP" command and then back to IDLE.

3. A trigger signal for executing a KR/IKR <ID> occurs (0703) :

10 The ProgramPointer (PP) is loaded with the address generated by LUT2. If the address is invalid, i.e. no entry is available for the KR/IKR to be loaded, it is loaded (0704) and the PP is reset.

15 The execution loop begins:

PP is incremented (this causes the BEGIN <ID> command to be skipped in the first loop pass), the occurrence of additional triggers is suppressed, RECONFIG is blocked. The commands are 20 executed and a jump is made to the beginning of the execution loop (0707) .

25 The "STOP" command is executed separately (0705) . The trigger and RECONFIG are enabled again and the state machine jumps to IDLE.

30 The "EXECUTE" command is also executed separately (0706) . The ID provided in EXECUTE <ID> is written into the ID REG. PP is reloaded and the KR/IKR provided via ID is executed (0708) .

After resetting the CT, the base configuration is loaded into the CTR and a jump is made directly to the execution of the basic configuration (0709) .

35 Figure 8 shows the structure of an SCRR ARB. The signals to be arbitrated go via DataIn to a mask (0801), which lets through or blocks a contiguous part of the signals according to the

known table. A conventional priority arbiter (0802) according to the related art arbitrates a signal from those let through and delivers its binary vector (BinaryOut) together with a valid/invalid identifier (ValidOut) (also according to the related art) as an output of the SCRR-ARB.

This signal is decoded according to the known table (0803) and sent to a register for clock synchronization (0804). The DataIn mask is sent via this register. The register is controlled either by a clock signal or by a Next signal (Enable EN) which queries the next valid binary vector. In the case of a reset or if the indication of the ValidOut identifier is invalid, the register is switched so that the DataIn mask lets all signals through.

The structure of the mask is shown in 0805. In 0806 the mask is illustrated again, signals DataIn 0 - DataIn 1 having the same priority according to the SCRR principle while DataIn m - DataIn n have a higher priority.

Figure 9 shows the LUT structure. The binary vector (BinaryIn) of the arbitrated trigger is supplied to the address input of LUT1 (0901). LUT1 translates the binary vector either into a valid trigger in order to forward the latter to another CT or into a valid ID. Both are output via 0910. 0911 shows whether this is a trigger or an ID.

If no translation of the incoming binary vector is entered in LUT1 via the "REFERENCE" command, the "Illegal Trigger" signal 0914 is generated via a bit entry or a comparator upon a certain token (e.g., "VOID").

A trigger is sent to external CTs via 0912; IDs are further processed via the multiplexer (0902). 0902 switches either the data output of LUT1, which provides a valid ID, or the ID register (0903) of the CT to the address input of LUT2 (0904). 0904 has a cache-like structure, i.e. the less significant

part (0906) of the data output of 0902 is switched to the address input of 0904, while the more significant part (0907) is switched to the data input of 0904. The data output belonging to 0907 is compared to 0907 by a comparator (0905).

5 The advantage of this procedure is that 0904 does not have to have the depth to translate all IDs, but may be considerably smaller. Like a conventional cache, only some of the IDs are translated, it being possible to establish in the LUT2 with the help of 0907 whether the selected ID corresponds to the
10 one specified by LUT1. This corresponds to a cache/TAG procedure according to the related art.

A multiplexer 0908 is assigned to a second data input of 0904; depending on the operation, this multiplexer provides the
15 FreePointer (FP, LOAD operation), the GarbagePointer (GP, REMOVE operation) or an invalid identifier/token (NoAdr, REMOVE operation) for storage at LUT2. Both pointers refer to memory locations in the CTR; "NoAdr" indicates that no entry exists for the appropriate ID; the entry has been deleted.
20 This is determined at the data output by comparing the data with the "NoAdr" token via comparator 0909. The following is forwarded to the state machine:

- the occurrence of a binary vector, via "ValidIn" (see
25 Figure 8).
- the information of whether a trigger or an ID results from the translation in LUT1 (0911, "Trigger/ID Out"). Triggers are forwarded to other CTs via 0912; IDs are processed in their
30 own CTs and forwarded to LUT2.
- the result of 0905 indicating whether the corresponding ID is stored in 0904 ("Hit/Miss Out").
- the result of 0909 indicating whether the respective ID points to a valid address in the CTR ("NoEntry Out").
35

The address generated by 0904 is forwarded to the CTR ("CTR Address Out").

5 LUT1 is loaded via the "REFERENCE" command with the translation of the incoming binary vector to a trigger or an ID. The operands of the command are sent to LUT1 via the 0913 bus. The ID register (0909) is loaded via the same bus.

10 Figure 10 shows the pointer arithmetics of the GarbagePointer (PG) [sic], ProgramPointer (PP), MovePointer (MP) and FreePointer (FP). Each pointer has a separately activatable loadable up/down counter. Each counter can be loaded with the value of each other counter, if necessary, as well as with the output of LUT2 (1007).

15 It is determined via the comparator whether

1. PP is equal to MP
- 20 2. MP is equal to FP
3. FP is equal to the maximum position in CTR. The results are used for controlling the state machine.

25 One of the pointers is sent to the CTR address input via a multiplexer (1001). The data arrives at the CTR via a multiplexer (1002) either from the higher level CT (1005) or from a register (1003). The data from the higher level CT or the CTR is forwarded to the state machine and the FILMO (1006) via a multiplexer (1004). If a REMOVE command occurs, the data is sent from the higher level CT directly to the state machine via 1004, while otherwise the commands are sent from the CTR to the state machine [German text garbled]. Register 1003 is used for storage and feedback of commands, shifted from one address to another during one pass of the Garbage Collector, to the CTR input.

The structure of a FILMO is illustrated in Figure 11. The data arrives from CTR (1101) in FILMO and it is either written into the FILMO RAM (1103) via the multiplexer (1102) or sent to the configurable elements (1116) via the multiplexer (1104). If
5 data is deleted in 1103, a "NOP" token is written after 1103 via 1102. The "NOP" token is recognized by the comparator (1105) at the data output and a write to the configurable elements is prevented. Either the write pointer WriteP (1107) or the read pointer (1108) is sent to the address input of
10 1103 via multiplexer 1106. The read pointer is saved in register 1109 in order to allow a reset (see Figure 6).

The "Full" fill status counter (1110) of 1103 is stored in register 1111 for resetting according to Figure 6. Two
15 comparators test whether 1103 is empty (1112) or full (1113). It is selected via multiplexer 1115 whether the control signals of the state machine (from 1101) or of the FILMO are sent to 1116.

20 Figure 12a shows the hierarchical structure of the CTs. All CTs obtain their data from the ROOT CT (1201) and the ECR assigned to it (1204). One or more CTs exist for each implementation level in a module. Each CT is responsible for managing its level and the lower-level CTs. It is not
25 necessary for all branches of the tree to have the same depth. For example, there may be fewer levels for controlling the periphery (1202) of a module than for controlling the operating units (1203). Data transfer takes place in the form of a tree. Each CT operates as a cache for all the lower-level
30 CTs.

Figure 12b shows the trigger flow between the CTs. While data flow takes place in the form of a tree, the trigger flow is undetermined. Each CT can send a trigger to each other CT. In
35 general, trigger exchange only takes place from the leaves (1203) in the direction of the ROOT CT (1201). Occasionally, however, transfer may also take place in the opposite

direction.

Figure 12c shows a trigger vector broadcast, with 1205 sending a trigger vector to all CTs.

5

Figure 12d shows a HIGHER trigger vector which 1206 sends to the higher level CT. 1207 sends a LOWER trigger vector to all the lower-level CTs. 1208 transmits a directly addressed (ADDRESSED) trigger vector to a certain CT which is not directly connected to 1207.

10

In Figure 13 two independent IKR n and m request a common KRx cached in the higher level CT. It is indicated that this KR is cached by the entire branch and is also available in an adjacent branch (1301) via a common CT.

15

Figure 14 shows an LUT system, modified with respect to Figure 9, used in ROOT CTs and CTs of the middle hierarchical levels. The basic difference with respect to the CTs described so far is that, instead of individual trigger signals, ID vectors and/or trigger vectors must be managed by the CT. A handshake signal (RDY) is assigned to each vector to indicate the validity of the vector, which are sent to an arbiter (1401). One of the trigger vectors (1404) or one of the ID vectors (1405) is selected via the multiplexers (1402, 1403). Trigger vectors go directly to the address input of LUT1 (1406), which is otherwise wired according to Figure 9. The ID register (1407) is also wired according to Figure 9. Contrary to Figure 9, multiplexer 1408 has three inputs (see 0902). The multiplexer is activated by arbiter 1401, in addition to the state machine. ID vectors are directly forwarded to LUT2 via 1403 over the additional input. Bus 1409 is used for this purpose. (In principle, IDs can be directly switched to LUT2 via a multiplexer (1408) even in the case of CTs according to Figure 9. The IDs can then be sent without translation directly from the CEL to the LUT2.) "Trigger/ID Out" is generated according to Figure 9. A "ValidIn" signal which is

forwarded onto a "Valid Out" according to Figure 9 does not exist. Instead, depending on the arbitration by 1401, a "Valid Trigger Out" is generated for trigger vectors and a "Valid ID Out" is generated for ID vectors in order to inform the state machine on how the processing is to take place.

Bus 1409 is connected to another unit via 1410; this unit only exists in ROOT CT and is described in Figure 15.

A ROOT CT requires, in addition to the normal CT functions, an interface to the external configuration memory (ECR), as well as the required address generator and units for managing access to the ECR.

A conventional CT translates the trigger vectors received in LUT1 to an ID and the ID in LUT2 to a memory location in the CTR (see Figure 16a). A ROOT CT translates an ID within the ECR, upon access to the ECR, to an address in the ECR where the KR/IKR referenced by ID begins. For this purpose, a memory range is established in the ECR, whose size corresponds to the possible number of IDs (for example, if an ID is 10 bits wide, $2^{10} = 1024$ possible IDs result, i.e., 1024 entries are reserved in the ECR). In the following examples, this memory range is located at the lower end of the ECR and is referred to as LUTECR in order to emphasize the similarity with LUT2. The translation of a trigger to an ID takes place according to the known CTs in LUT1 (1601). For greater clarity, Figure 16b illustrates access to the ECR.

In Figure 15 an ID goes to multiplexer 1501 via bus 1410 of Figure 14. The ID is written into loadable counter 1502 via 1501. The output of 1502 goes, via a multiplexer 1503, to the address bus (1504) of the ECR. The translation of the ID into a memory address goes via databus 1505 through a multiplexer/demultiplexer (1506) to 1501, which loads 1502 with the memory address. Subsequently the data words of the corresponding KR/IKR are read from the ECR via the LOAD ECR

state machine (see Figure 17) and written into the CTR, 1502 being incremented after each read sequence, until the "STOP" command is read.

5 The higher level HOST writes KR/IKR into the ECR via 1503/1506 through interface 1507. The state machine (CTS) arbitrates whether the HOST or the ROOT CT has access to the ECR.

10 After resetting the module, a base configuration (BOOT KR) must be loaded. For this purpose, a fixed memory address (BOOT ADR) is introduced which points to the first memory location of the BOOT KR. The memory location 0h is recommended as the BOOT ADR if the IDs start with 1; otherwise 2^{ID} or some other
15 memory location can be used. In the exemplary embodiment, 2^{ID} is used.

20 The ROOT CT performs a lookup in order to load the BOOT KR at the location BOOT ADR if a BOOT KR is loaded. The ROOT CT
writes the data after 1502 in order to load the BOOT KR from there until a "STOP" command occurs.

25 A monitoring unit within the ROOT CT assumes the synchronization of the HOST with the module. This takes place as follows:

30 The addresses small 2^{ID} [sic] are monitored by 1508, i.e., when the HOST accesses these addresses, a signal (ACC ID) is sent to the state machine (CTS).

35 BOOT ADR is also monitored via 1509 and sends an ACC BOOT signal to the state machine (CTS).

The state machine (CTS) responds as follows:

- if HOST writes to the BOOT ADR, this causes BOOT KR to load.

- if HOST writes the data word 0 (1512) to the BOOT ADR, this is confirmed via comparator 1510 and causes the module to stop.

5 - if the HOST writes to a small 2^{ID} address, the address is loaded into the REMOVE register (1511). Since the address corresponds to the ID (see ECR LUT), the ID of the modified KR/IKR is positioned in 1511. The REMOVE <ID> command for immediate execution is sent to all CTs (1513). The CTs then
10 delete the KR/IKR of the corresponding ID from their CTR, i.e., LUT2. Upon the subsequent call up of the KR/IKR, the CTs must necessarily load the new KR/IKR from the ECR.

15 Figure 17 shows the sequence in a ROOT CT when a KR/IKR is loaded from the ECR. If an ID is not found in the internal CTR (see Figure 1, 1701), the ID is written into counter 1502 (1703). An access to the ECR having the address in 1502 delivers the base address of the KR/IKR. This is written into 1502 (1704). A LOAD according to Figure 2 takes place (1702).
20 The data is read from the ECR (1705) instead of a higher level CT and is not only written into the CTR of the same CT, but is also sent to the lower-level CT (1706).

25 In a mid-hierarchical level CT, the translation of the triggers takes place as in Figure 1, with the exception that trigger vectors and ID vectors are handled according to Figure 14. The KR/IKR are loaded according to Figure 2, with the exception that the data words are not only written into the CTR of the same CT (0210), but are also sent at the same time
30 to the lower-level CT.

35 Figure 19 illustrates the FILMO principle. The FILMO (1901) is always run through from the beginning to the end during read and write accesses (1902). If entries are written and deleted from the beginning of the FILMO (1903), the read pointer is shifted to the first undeleted entry (1904). If entries are written from the middle of the FILMO (1905), the read pointer

remains unchanged (1906), and the entries are marked with "NOP" (1907). If data is written into the FILMO (1908), it is appended to the end of it after the last entry (1909). The read pointer (1910) remains unchanged.

5

Of course, a CT may also be designed with a single memory which includes LUT1, LUT2, and CTR. However, control is more complex in this case. The CTs have a design similar to that of ROOT CT, which already integrates LUT2 and the CTR in the ECR.
10 A description of these CTs is unnecessary for understanding the procedure.

10

If a CT is used as a cache system for data, triggers are introduced for writing data into the CTR. The data is written
15 from a CEL into the CTR. The necessary changes are trivial; the FILMO is no longer required.

15

In caching data, the problem of data consistency arises. It can be circumvented by using a method according to German Patent Application 42 21 278 A1 in order to identify the data and its validity at the individual hierarchical levels. If data is requested for executing a read-modify-write cycle (RMW cycle), the data is marked at all hierarchical levels as INVALID using an additional entry in the CTR/ECR. The unique ID of the KR/IKR using the data can be inserted in the entry for this purpose. The data cannot be used by any KR/IKR having another ID until the KR/IKR using the data has written back the data (see write-back method according to the related art) and deleted its ID.

25

30

Figure 20 shows one embodiment:

35

In Figure 20a CT 2007 requests data from the higher level CT, which in turn requests data from ROOT CT 2004; the ID of the requesting KR/IKR (2001) is transmitted with the data request. The data (2002) is sent to 2007. All other subsequent accesses are rejected (2003).

In Figure 20b the data is written back (2005); other subsequent accesses are accepted again (2006).

5 In Figure 20c data is requested from a mid-hierarchical level CT, which has the data and sends its to 2007. The ID for blocking the data is sent to all CTs in that hierarchy (2001). When the data is written back (Write-Back) in Figure 20d, the data is written to all CTs of that hierarchy and the ID is deleted.

10

Figure 21 shows the communication of an INITIATOR CT (2101) via a plurality of intermediate CTs (2104, 2105, 2106) with a TARGET CT (2102), as well as direct communication without the intermediary levels with a TARGET CT (2103) by the
15 GETCTS/LOOSECTS method.

20 2101 establishes a link to 2103. After successfully establishing the link, 2101 receives a GRANT from 2103 as an acknowledgment. Subsequently 2101 establishes the link to 2102 via 2104, 2105, 2106. The link to 2102 is not acknowledged (GRANT) until 2102 is reached.

25 If the link cannot be established because one of the buses is busy, a REJECT is sent to 2101, and 2101 interrupts the sequence. This means that the link to 2103 is also interrupted and a REJECT is sent to 2103.

30 However, if 2102 acknowledges the link with GRANT, 2101 sends a GO command to 2103 and 2102 to acknowledge the successful establishment of the bus link and the synchronization to 2103 and 2102 at the same time. Data or commands can be transmitted through this protocol synchronously and without deadlocks, since GO ensures that all TARGETs correctly receive the commands.

35

Figure 22 shows the sequence of the inter-CT communication by the GETBUS/LOOSEBUS method. While in the method according to

Figure 21 the higher level CTs have the controlling and prioritizing task, control is assumed here by the inter-CT buses (2201).

5 A link to 2103 is established by the fact that the INITIATOR CT (2101) requests its local inter-CT bus (2202). Requests are acknowledged if the bus is free (ACCEPT) or rejected if the bus is busy (REJECT). Subsequently it sends the address of 2102 to the bus. According to the addressing plan, the bus
10 control recognizes that the address is outside the local bus addresses and establishes a link to its local bus (2203) via the higher level CT 2104. As the address of 2102 is within its address range, the link to the local bus is established by 2102 (2204) via 2106. Since 2101 is now the sole bus master of
15 all the buses required for the data communication, it is ensured that problem-free and deadlock-free communication takes place, since the communication channels are blocked for all the other CTs. 2102 and 2103 also cannot use the buses, since, in their TARGET role, they can only receive commands
20 and can only send data themselves upon request by the INITIATOR (2101).

As soon as the communication is terminated, the buses are freed by a signal from 2101.

25 If 2101 encounters a busy bus during the establishment of the link, a REJECT is sent to 2101, and 2101 frees the bus systems again and attempts to establish the link at a later time. If multiple CTs request the same bus at the same time, the higher
30 level CT has a higher priority (2205). This prevents a link that has reached an advanced stage in being established and already runs over multiple levels from being interrupted by a still very local bus buildup.

35 An extended protocol makes it possible in the case of a REJECT to free only those buses that are required by the bus buildup having a higher priority. This can considerably enhance

performance, since not all links have to be established anew at a later time.

The structure of the inter-CT bus for the method according to
5 Figure 22 is shown in Figure 23. CTs 2301-2304 are connected, together with higher level CT 2305 (interface 2307), to inter-CT bus 2312 via their interfaces (2308-2311). Connection to the inter-CT bus takes place via a round-robin arbiter which provides 2308-2311 with the same priority and 2307 with a
10 higher priority; it activates a multiplexer to connect the buses (2306). A state machine which evaluates the control signals (e.g. establish/remove links, ACCEPT, REJECT, etc.) is assigned to the arbiter.

15 Figure 24 shows the structure of the address plan of a unidimensional CT tree. The rectangles symbolize a CT. The address of the CT has been entered. "--" identifies irrelevant address bits, which are not evaluated; the relevant address bits are given as binary 0 or 1; "*" stands for any desired
20 address bit. It can be easily seen that this design can also be used for multidimensional trees by projection in which case each of the given addresses represents one of the axes; in other words, a corresponding separate address system exists for each axis.

25 Figure 24a shows the addressing from CT 0001. The relative address -1 is specified. By calculating $-1+1 = 00$ ("relative motion" + "address of the INITIATOR CT on current level"), the CT 0000 connected to the same local bus can be calculated.

30 In Figure 24b, CT 0010 calls up the relative address +10. The calculation $10+0 = 10$ ("relative motion" + "address of the INITIATOR CT on current level") provides the carryover 1, since the address range of the lowermost local bus is exactly
35 one bit wide. Therefore the next higher bus is selected. This address calculation with $10 + 10 = 100$ ("relative motion" + "address of INITIATOR CT on current level") again provides a

carryover, since its address range (2 bits) is exactly one bit larger than the address range of the lowermost bus. On the next higher level no carryover occurs with the calculation $10+010 = 0100$ so that the third bit (from the left) addresses path 1^{**} having the next lower level; the second bit (from the left) addresses path 10^* of the next lower level and, finally, the last bit selects the TARGET CT.

Figure 24c shows the known method over two levels in the positive direction, and Figure 24d shows the method over three levels in the negative direction with a negative overrun.

Figure 25 shows the structure of a two-dimensional CT tree. The CTs (2501) are located on the lowermost level (2502) in a two-dimensional arrangement. The address of the dimension is provided in the respective CT with x/y. The next higher level (2504) is above 2502. Each of its CTs (2503) controls a group of four CTs of the level 2502. The address space of the CTs at 2504 is one bit wider; * identifies the address bits of level 2502 which are irrelevant for the selection of the CTs at 2504. ROOT CT 2505 is one level above 2504. Its address is again one bit wider; the meaning of * is equivalent

Figure 26 shows the linkage of the garbage collector in the case of the microcontroller implementation. All KR_s are linked to one another via the header entries (garbage-previous/garbage-next). When the garbage collector runs through the list, the age of the KR is documented for the cache statistics (2602) by incrementing the entry by (+1). The garbage collector looks for the KR statistics entry (2601) which shows whether the KR still exists in the FILMO list. In this case KR should not be deleted by the GC since it still contains unconfigured KWS. As an alternative, this test could also run via the FILMO-next and FILMO-previous entries.

Figure 27 shows the linkages of the FILMO list.
The linkage can be completely different from that in the

garbage list (Figure 26). The KRs are linked via FILMO-previous and FILMO-next. The KR statistics (2701) entries point to the first not yet configured KW in the respective KR. A FILMO run is structured so that the KR is started in the 5 first ID. After execution, the position of the unexecuted KW is written after 2701. If KR has been completely executed, the KR is removed from the linked FILMO list, but remains in the memory. Then a jump is made over the FILMO list, to the next KR which is also processed.

10

Figure 28 illustrates the structure of a KR in the case of microcontroller control. At the beginning there is a jump command, which jumps behind the header (2801) of the KR. The FILMO bit (2802) is assigned to each KW. A 1 (2803) shows that 15 the KW has been accepted (ACCEPT) by the CEL and is no longer executed in the next run. A 0 (2804) shows a REJECT; the KW must be executed again in the next run. The optional KR statistics (2701) point to the first KW marked with 0. If PUSHRET (2805) receives a REJECT, the processing of the KR is 20 interrupted here and set up again in the next run either at the first KW or to the location at which 2701 points. Otherwise the KR is terminated properly at its end 2806.

Figure 29 shows the circuit for saving the status information 25 of a CEL prior to running through the FILMO or starting a KR. The status information goes from the CEL (2901) to a register (2902). Prior to running through the FILMO or starting a KR, the CT sends an enable signal (2903) to 2902. Then the status information is accepted and forwarded to the CT (2904). 2904 30 remains constant until the next transmission from 2903.

Definition of Terms

ACCEPT Signal showing that the addressed CEL is in a configurable state and accepts the transmitted KW.

5

Block commands (such as BLOCK-MOVE) Commands that shift a plurality of data (a block) in the memory or between the memory and the periphery. The origination address of the data to be shifted, the target address of the data, and the length of the data block are provided.

10

Broadcast Sending of information to a plurality of receivers.

15 Data receiver The units(s) that further process(es) the results of the CEL.

Data transmitter The units(s) that make(s) data available to the CEL as operands.

20

Data word A data word has a bit array of any desired length. This bit array represents a processing unit for a system. Both commands for processors or similar modules and pure data can be encoded in a data word.

25

Deadlock State in which no data processing is possible due to mutual blockage.

30 DFP Data flow processor according to German Patent/Offenlegungsschrift 44 16 881.

DPGA Dynamically configurable FPGAs. Related art.

35 Elements Umbrella term for all types of standalone units that may be used as pieces in an electronic module. Thus, elements include:

- configurable cells of all types

- clusters
- RAM blocks
- logic
- arithmetic units
- 5 - registers
- multiplexers
- I/O pins of a chip

Event An event can be evaluated by a hardware element in
10 any manner that is fitting for the application and can trigger
a conditional action as a response to this evaluation.

Thus, events include, for example:

- clock cycle of a computer system
- internal or external interrupt signal
- 15 - trigger signal from other elements within the module
- comparison of a data flow and/or a command flow with a value
- input/output events
- expiration, overrun, reset, etc. a counter
- evaluation of a comparison.

20

FIFO First-in first-out memory according to the related
art.

25 FILMO Modified FIFO from which linear data is read. The
read pointer is not restricted to the beginning of the memory.

FPGA Programmable logic module. Related art

30 F-PLUREG Register in which the function of the CEL is set.
The OneSHOT and Sleep modes are also set. PLU writes into the
register.

35 Fragmentation Fragmentation of the memory into a plurality of
often small and therefore useless memory ranges.

Garbage collector Unit for managing the memory. Prevents
fragmentation.

H level Logical 1 level depending on the technology used.

HOST A computer on a level higher than the module or component.

5

IDLE cycle The cycle in which a state machine performs no processing. Basic state of a state machine.

INTER-CT BUS Bus system between the CTs of one level and a
10 higher level CT (or CT group).

INITIATOR CT which initiates access to an inter-CT bus.

Pointer pointer to an address or a data word.

15

Configurable element (KE) A configurable element represents a unit of a logic module which can be set for a special function via a configuration word. Configurable elements are therefore all types of RAM cells, multiplexers,
20 arithmetic logical units, registers, and all types of internal and external interconnection descriptions, etc.).

Configurable cells (CEL) See logic cells.

25 Configuring Setting the function and interconnection of a logical unit, an (FPGA) cell, or a CEL (see Reconfiguring).

Configuration data Any number of configuration words.

30 Configuration routine (KR) Multiple configuration words combined to form an algorithm.

Configuration memory The configuration memory contains one or more configuration words.

35

Configuration word (KW) A configuration word has a bit array of any desired length. This bit array represents a valid

setting for the element to be configured so that a functional unit is obtained.

5 Load logic Unit for configuring and reconfiguring the CEL.
Structured by a special microcontroller adapted to its function.

10 Logic cells Configurable cells used in DFPs, FPGAs, DPGAs, which perform simple logical or arithmetic operations according to their configuration

Lookup table Known method for translating data.

15 LUT1 Lookup table which translates a trigger to an ID and determines whether the trigger is assigned to a valid ID.

20 LUT2 Lookup table which translates an ID to the address of the respective KR in the local memory and determines whether the KR is available in the local memory.

L level Logical 0 level, depending on the technology used.

25 Mask Bit combination providing the valid signals within a plurality of signals.

Prioritizing Determining a sequence order.

RECONFIG Reconfigurable state of a CEL

30 RECONFIG trigger Setting a CEL into the reconfigurable state.

REJECT Signal which shows that the addressed CEL is in a non-configurable state and does not accept the KW sent.

35 REMOVE <ID> 1. Command within a KR to remove the KR referenced by its ID.

2. Command of a higher level CT via a separate interface or handshaking to a lower level CT to delete the KR referenced by its ID.

5 RESET Resetting a module or an entire computer system into a certain basic state.

ROOT CT CT of the highest hierarchical level having direct access to the external configuration memory.

10 Round-robin arbiter Arbiter running in a circle that always assigns the lowest priority to the last-arbitrated signal.

15 State machine

Sync signals Status signals generated by a configurable element or an arithmetic unit and forwarded to other configurable elements or arithmetic units to control and 20 synchronize data processing. A sync signal can also be sent back to the same configurable element or arithmetic unit with a time delay (stored).

25 TARGET CT which is a target of an access to the inter-CT bus.

Trigger Synonym for sync signals.

30 Reconfiguring Reconfiguring any desired number of CELs while any number of other CELs continue to perform their own functions (see configuring).

Linked list Data structure linked together via pointers according to the related art.

35 Cells Synonym for configurable elements.

State machine Logic capable of assuming different states. The transitions between the states depend on different input parameters. These machines are used for controlling complex functions and are known from the related art.